CLAIMS

1. A data rate controller for controlling a rate that data is transferred over a backplane in a network processing device, comprising:

a bandwidth allocator configured to allocate bandwidth to an input port for transmitting data over the backplane to an output port;

a bandwidth limiter configured to identify a maximum allowable bandwidth the input port is allocated on the backplane; and

a bandwidth tracker configured to identify an amount of bandwidth currently allocated to the input port for transmitting data over the backplane to the output port, the input port prevented from connecting to the output port when the current allocated bandwidth is used up.

2. A data rate controller according to claim 1 wherein the bandwidth allocator comprises a register that stores a programmable peak time slot rate value.

3. A data rate controller according to claim 2 wherein the bandwidth allocator comprises a counter that assigns an additional amount of bandwidth when the counter decrements the peak time slot rate value down to zero.

4. A data rate controller according to claim 1 wherein the bandwidth tracker includes a counter that decrements the amount of bandwidth currently allocated when the input port is connected through the backplane to an output port and increments the amount of currently allocated bandwidth when the input port is not connected through the backplane to the output port.

5. A data rate controller according to claim 4 wherein the bandwidth tracker is disabled from counting up when the maximum allowable bandwidth has been reached.

6. A data rate controller according to claim 1 including a network processing device that includes multiple input ports and multiple output ports and

including one data rate controller for each input-output port combination in the network processing device.

7. A data rate controller according to claim 6 wherein the bandwidth tracker prevents requests to the arbitration circuit for the input ports having exhausted bandwidth allocation.

8. A data rate controller according to claim 6 wherein each one of the input ports has associated virtual output queues each associated with a different one of the output ports and the virtual output queues each have associated data rate controllers.

9. A data rate controller according to claim 1 including an arbitration circuit configured to arbitrate between input ports for connections to output ports, the arbitration circuit selecting the input ports for a next time slot according to both a priority and weight of packets at the input ports.

10. A data rate controller according to claim 9 wherein the arbitration circuit selects the input ports in a round robin order when two or more of the input ports have a same highest priority and a same largest weight.

11. A data rate controller according to claim 8 wherein the arbitration circuit conducts output port arbitrations for all of the virtual output queues dedicated to the same output ports and conducts input port arbitrations between the virtual output queues for the same input port issued grants during the output port arbitrations, the data rate controllers associated with the virtual output queues preventing participation in the output port arbitrations when bandwidth allocation associated with the virtual output queues is exhausted.

12. A method for controlling a rate that data is transferred over a switch fabric, comprising:

allocating bandwidth to input ports for transferring data to output ports over the switch fabric;

sending requests from the input ports for connecting to the output ports during a next time slot;

increasing bandwidth allocation for the input ports that are not connected to the output ports for the next time slot;

decreasing bandwidth allocation for the input ports that are connected to the output ports for the next time slot; and

preventing the input ports from sending requests for the input ports when the bandwidth allocated to the input ports has been exhausted.

13. A method according to claim 12 including identifying maximum bandwidth allocations for the input ports and disabling bandwidth allocation to the input ports that reach these maximum allowable bandwidth allocations.

14. A method according to claim 12 including providing a programmable amount of bandwidth allocation for the input ports.

15. A method according to claim 12 including:

selecting a bandwidth allocation to time slot period ratio; and

assigning a bandwidth allocation value to the input ports according to the selected ratio.

16. A method according to claim 12 including:

identifying bandwidth allocated to the output ports for transferring data to a network; and

preventing the output ports that have used up this allocated bandwidth from granting connections to the input ports.

17. A method according to claim 12 including:

providing multiple input port buffers for each input port, each input port buffer associated with a different output port;

sending connection requests from multiple ones of the input port buffers for connecting to the output ports through a switch fabric;

conducting arbitrations for the connection requests to determine which input ports are connected to which output ports; and

disabling the input port buffers from sending connection requests when their allocated bandwidth has been used up.

18. A network processing device, comprising:

multiple input ports for receiving incoming packets;

multiple output ports for outputting packets;

a switch fabric coupled to the different input ports and the different output ports;

multiple virtual output queues associated with each one of the input ports, each one of the virtual output queues dedicated to a different one of the output ports; and

a set of data rate controllers associated with each one of the virtual output queues for controlling a data rate that the input ports can transfer data to the output ports over the switch fabric.

19. A network processing device according to claim 18 including a scheduler that configures the switch fabric for connecting the input ports to the output ports.

20. A network processing device according to claim 19 wherein the data rate controllers prevent input ports that have exceeded a data rate limit from sending connection requests to the scheduler.

21. A network processing device according to claim 20 wherein packets are transferred over the switch fabric during time slots and the data rate controllers control data rate per time slot ratios between the input ports and output ports.

22. A network processing device according to claim 21 including a second set of data rate controllers that control a rate that data is received by the output ports from multiple input ports.